

*File*

DRAFT:

~~5/19/67~~

6/19/67

*? published*

Identification

Hard-Core Supervisor Callout Mechanism

R. L. Rappaport

Purpose

At system initialization time, the intersegment references made between hardcore supervisor modules are prelinked. References from hardcore modules to segments in other rings cannot be linked at initialization time since the segments to which the references are directed, are not available at this time. Linking at execution time cannot occur in the normal fashion (i.e. by an ftz fault in the hardcore ring) since the linker is in the administrative ring. The mechanism described here allows the linking to be accomplished at execution time by arranging for the linkage faults to occur in the ring to which the call out is directed.

*read  
not all  
outs from  
ring 0*

Introduction

In order to simplify linking of calls out from the hardcore ring all the calls are directed through one procedure. That is, only one hardcore module ever calls out to any segments not in the hardcore ring. This procedure is known as hcs\_out and it has an entry point for each distinct procedure which can be called from a hardcore ring module. Calls out are then handled in the following way.

Hcs\_out has a distinct copy of its linkage section in each protection ring. Hcs\_out is always entered through the copy of its linkage section in the hardcore ring. Upon entry, hcs\_out determines to which protection ring the call is directed, and ~~then~~ <sup>reads</sup> base pair lb ← lp to point to the copy of the linkage

*Answer*

section that resides in that ring. At this point, hcs\_out merely calls, that is, transfers through its new linkage section<sup>to</sup> the desired procedure and the linkage fault will occur in this new linkage section residing in the entered ring.

Let us present an example to clarify the above. Figure 1 illustrates the problem.

Suppose segment <a> in the hardcore ring desires to call segment <b> in the administrative ring. The code in segment <a> would be:

```
call hcs_out $ b;
```

At system initialization time <a.link> is linked directly to <hcs\_out.link>, the linkage section of hcs\_out in the hardcore ring. The flow of control is as follows: <a> transfers through its linkage section to the hardcore linkage section of hcs\_out, which in turn transfers directly to the entry hcs\_out \$ b. Hcs\_out knows<sup>(or can calculate)</sup> that <b> resides in the administrative ring, ~~and~~

Therefore hcs\_out executes an eapl for <hcs\_out1.link> and then merely calls <b>, that is, transfers through <sup>hcs\_out1.link</sup> <~~hcs\_out.link~~> to <b.link>.

If this is the first time that <b> has been called from hcs\_out, then a linkage fault will occur in <sup>hcs\_out1.link</sup> <~~hcs\_out.link~~>.

### Discussion

Each process has its own distinct version of hcs\_out. Segment <hcs\_out> consists of a number of entries and an array of pointers to its various linkage sections, one per ring. Initially this array of linkage pointers is filled with "software simulated" fault word pairs (see Section BB.5.03).

Reference to one of these fault word pairs will generate software simulated fault number 4. This fault is serviced in the hardcore ring by creating a copy of  $\langle \text{hcs\_out} \cdot \text{link} \rangle$  in the appropriate ring and filling the faulting word pair with a pointer to this new linkage section. In this way the various linkage sections can be created dynamically only as they are needed.

Hcs\_out is used by all procedures, pre-linked at system initialization time, which have reason to call procedures in other rings. For example, the fault interceptor, which resides in all rings, is pre-linked at initialization time. All references by the fault interceptor to procedures outside the hardcore ring are therefore directed through hcs\_out, although if the fault interceptor is in ring n calling a ring n procedure, one would not ordinarily consider this a call out from the hardcore ring.

Let us go through an example to illustrate this. Suppose the fault interceptor operating in ring 3 wants to call subroutine <sup>in ring 3</sup> signal. The flow of control is as follows (see figure 2):

1. The fault interceptor ( $\langle \text{fim} \rangle$  in the diagram) executes a call that looks like:

call hcs\_out\$signal (arglist)  
 where arglist is the list of arguments expected by the called procedure. This  
 results in a transfer through its linkage section ( $\langle \text{fim} \cdot \text{link} \rangle$ )  
 to the linkage section of hcs\_out, with which it was pre-linked  
 at initialization time. That is,  $\langle \text{fim} \cdot \text{link} \rangle$  transfers to  
 $\langle \text{hcs\_out} \cdot \text{link} \rangle$ . This implies that both  $\langle \text{fim} \cdot \text{link} \rangle$  and  
 $\langle \text{hcs\_out} \cdot \text{link} \rangle$  reside in each ring.

2. In  $\langle \text{hcs\_out} \cdot \text{link} \rangle$  is a standard entry which sets lb ← lp to point to this linkage section and transfers to the procedure. That is, even though this is being executed in ring 3, entry to hcs-out is still made through its hardcore linkage section.

3. Hcs\_out then calls subroutine getring to find out what ring it is currently executing in. That is, it ~~calculates~~ determines to which ~~linkage section~~ linkage section it should switch.

4. Hcs\_out then reloads lb ← lp to point to the copy of its linkage that exists in the ring currently active. [Reloading the linkage pointer ~~can~~ will cause generation of software simulated fault number

4 if this ~~is~~ is the first time hcs\_out has executed in this ring. This potential fault is irrelevant to the operation of hcs\_out. Its handling makes it completely transparent. The handling of this fault is described below. ]

5. Finally, hcs\_out merely calls signal. That is, it transfers through its linkage section to signal possibly causing a linkage fault to occur. Entry to signal is made through its linkage section in that ring.

[Note: Segments which reside in more than one ring must have resident linkage segments in these rings. If the segment expects to dynamically link to other segments, as in the case of signal, then distinct copies of the linkage section must be in each distinct ring. If the segment is pre-linked and if its pre-linked linkage section is read only, as in the case of the fault interceptor, then the same linkage can reside in each ring.]

In this way one can see how the fault interceptor, which is linked to hcs\_out at system initialization time, is able to reach signal dynamically at execution time.

#### Software Simulated Fault Four

This fault is signaled in the hardcore ring. The faulting ~~and~~ ~~is~~ looks ~~is~~ is generated by the following; ~~set~~

eapl X, \*

X:

(77777)8	its
4	*

Any eapl instruction with an operand like this will generate the fault. The only legitimate place for such a ~~word pair~~<sup>set of words</sup> to exist is in the ~~array of~~

~~pointer~~ segment <hcs\_out>. The handler for this fault

(master mode since it must write in read only <hcs\_out>) checks to see if the ~~faulting word pair~~<sup>they</sup> do reside in this segment. If ~~they~~<sup>they</sup> do the appropriate

copy of <hcs\_out.link> is made in the appropriate ring and the word pair ~~at~~<sup>instruction</sup> X is replaced by an ITS pointer. If the faulting ~~pair word~~<sup>instruction</sup> is not in this

segment an error is signaled.

Hardware Ring

Administrative Ring




Figure 1. Numerals indicate sequential steps in the flow of control.




Figure 2.